



INVESTOR IN PEOPLE

The Patent Office
Concept House
Cardiff Road
Newport
South Wales
NP10 8QQ



I, the undersigned, being an officer duly authorised in accordance with Section 74(1) and (4) of the Deregulation & Contracting Out Act 1994, to sign and issue certificates on behalf of the Comptroller-General, hereby certify that annexed hereto is a true copy of the documents as originally filed in connection with the patent application identified therein.

In accordance with the Patents (Companies Re-registration) Rules 1982, if a company named in this certificate and any accompanying documents has re-registered under the Companies Act 1980 with the same name as that with which it was registered immediately before re-registration save for the substitution as, or inclusion as, the last part of the name of the words "public limited company" or their equivalents in Welsh, references to the name of the company in this certificate and any accompanying documents shall be treated as references to the name with which it is so re-registered.

In accordance with the rules, the words "public limited company" may be replaced by p.l.c., plc, P.L.C. or PLC.

Re-registration under the Companies Act does not constitute a new legal entity but merely subjects the company to certain additional company law rules.



Signed

Dated

19 FEB 2801

THIS PAGE BLANK (USPTO)


The Patent Office

1/77

Patents Act 1977
Rule 16

19MAY00 E538005-1 000611
P01/7700 0.00-0011978.4

Request for grant of a patent



The Patent Office

Concept House
Cardiff Road
Newport
South Wales NP10 8QQ

1.	Your reference	GB000009GB1		
2.	Patent application number (The Patent Office will fill in this part)	0011978.4		19 MAY 2000
3.	Full name, address and postcode of the or of each applicant (underline all surnames)	INTERNATIONAL BUSINESS MACHINES CORPORATION Armonk New York 10504 United States of America		
	Patents ADP number (if you know it)	519637001		
	If the applicant is a corporate body, give the country/state of its incorporation	State of New York United States of America		
4.	Title of the invention	REMOTE TRACING OF DATA PROCESSING NODES IN AN ASYNCHRONOUS MESSAGING NETWORK		
5.	Name of your agent (if you have one)	R D Moss		
	"Address for Service" in the United Kingdom to which all correspondence should be sent (including the postcode)	IBM United Kingdom Limited Intellectual Property Department Hursley Park Winchester Hampshire SO21 2JN		
	Patents ADP number (if you know it)	919005		
6.	If you are declaring priority from one or more earlier patent applications, give the country and the date of filing of the or of each of these earlier applications and (if you know it) the or each application number	Country	Priority App No (if you know it)	Date of filing (day/month/year)
7.	If this application is divided or otherwise derived from an earlier UK application, give the number and the filing date of the earlier application	No of earlier application	Date of filing (day/month/year)	

Is a statement of inventorship and of right to grant of a patent required in support of this request? (Answer 'Yes' if:

Yes

- a) any applicant named in part 3 is not an inventor, or
- b) there is an inventor who is not named as an applicant, or
- c) any named applicant is a corporate body.)

9. Enter the number of sheets for any of the following items you are filing with this form. Do not count copies of the same document

Continuation sheets of this form

Description 12

Claim(s) 3

Abstract 1

Drawing(s) 7

10. If you are also filing any of the following, state how many against each item.

Priority documents

Translations of priority documents

Statement of inventorship and right to grant of a patent (Patents Form 7/77)

Request for preliminary examination and search (Patents Form 9/77)

Request for substantive examination (Patents Form 10/77)

Any other documents (please specify)

11. I/We request the grant of a patent on the basis of this application

R D Moss

Signature

R D MOSS

18 May 2000
Date

12. Name and daytime telephone number of person to contact in the United Kingdom A Sekar
01962 818169

REMOTE TRACING OF DATA PROCESSING NODES
IN AN ASYNCHRONOUS MESSAGING NETWORK

Technical field of the Invention

The present invention relates to the remote tracing of a plurality of data processing nodes, which are connected to each other via a network.

Background of the Invention

As computer systems increase in size and complexity, it has become common to distribute data processing applications over a plurality of data processing nodes, communicating over a data communications network, for example, the Internet. This allows a large data processing task to be distributed over several data processing nodes whilst also allowing exchange of data messages, which can consist of requests for processing or replies containing the results of processing. The data processing nodes can be arranged so that a local data processing node can control a remote data processing node. However, it is difficult to monitor operation of remote nodes to facilitate error detection and correction.

Tracing operations are known in the prior art for the purpose of tracing the path of execution of an application to assist in locating errors therein. The tracing operation assists in problem determination by providing a snapshot record in storage of certain types of states existing when a location in an application is reached by the data processing node that is running the application. Such events or states are often stored in a trace table or in trace files in memory.

Static or local tracing operations in response to commands issued within a node is known in the prior art. However, these commands cannot be used to initiate remote traces or dynamic alteration of commands at runtime. Extending tracing operations to remote nodes may involve large overheads and is often intrusive. It is difficult to achieve an acceptable trade off between response times and existing hardware, as the link between data processing nodes may be slow, limiting the efficiency of data exchange. Additionally, if the results from the remote operations are written into a file for subsequent perusal or if the remote data processing node only periodically reports to the local data processing node, this will further delay response times.

Tracing operations are often implemented prior to debugging operations in order to isolate a problem area. Debuggers are tightly coupled to the processes which are targeted and are therefore effective in dealing in detail with the problem identified by the tracing operation. Debugging generally involves stepping a program one step at a time, through all possible paths of execution and monitoring its behaviour. This differs from monitoring the normal execution of the program in a production mode, when execution events are traced in real-time.

Remote debugging to facilitate problem determination and replication of real-time conditions on remote systems is known, for example, as described in US patent No. 5630049. Although avoiding problems of prior remote debuggers arising from the fact that the data processing nodes typically must be in close physical proximity to each other, which limits the flexibility of the testing environment, the debugging process carries a large overhead and is relatively intrusive in comparison with a trace process. Problems with data processing nodes on a customer site often occur in unpredictable circumstances and the fine control and detailed activity of a debugger is not required for first pass problem analysis.

US patent No. 5630049 also employs a method of asynchronous messaging in which messages between nodes of a network may be transmitted using process-private interrupts known as Asynchronous System Traps (ASTs). ASTs enable a faster and more reliable communication than with general asynchronous messaging.

Disclosure of the Invention

However, there is still a need for a remotely controlled tracing operation in a network of data processing nodes, which executes with minimum intrusion into the data processing nodes. There is also a need for the remote tracing operation to respond dynamically in an acceptable time frame and without limiting the physical proximity of nodes from each other.

Accordingly, the present invention provides a method for remote tracing from a local data processing node of the execution of a process within an application program running on a remote data processing node in a distributed data processing network, said application program including its own local trace facility, said nodes communicating by asynchronous messaging via a data exchange means and each node including process-private

interrupt handling means for indicating the presence of a command for the respective process in said data exchange means, said method comprising the steps of sending a trace command from a trace process running on said local data processing node into a data exchange means of said remote data processing node; in response to said trace command, causing a process-private interrupt of a target process running on said remote data processing node; in response to said process-private interrupt, said target process writing trace information from said trace facility to said data exchange means; transmitting said trace information across said network; receiving in a data exchange means on said local data processing node, said transmitted trace information; in response to receiving said trace information, causing a process-private interrupt of said trace process; and in response to said process-private interrupt, reading said trace information by said trace process, from said local data exchange means.

Specifically, the tracing operation is advantageous as it can be executed with networked data processing nodes. This allows an end user on a local data processing node to dynamically perform tracing and diagnostic operations of remote data processing nodes. For example, a service provider can perform online diagnostics of data processing nodes located on a customer site, assuming a secure architecture.

In a further preferred aspect of the present invention, there is provided a distributed data processing system comprising a plurality of data processing nodes connected via a network, each node having a processor, memory and operating system capable of executing application programs, each of said operating systems including data exchange means and interrupt handling means, a first of said nodes comprising means for sending a trace command from a trace process running on said first data processing node into the data exchange means of a second data processing node, said second data processing node including a trace facility for tracing the execution of a process within an application program running on said second node, said interrupt handling means of said second data processing node in response to said trace command, causing a process-private interrupt of a target process running on said second data processing node; said second data processing node further including means for writing trace information from the trace facility, to said second data exchange means in response to the process-private interrupt; and means for transmitting said trace information across said network whereby said first data exchange means receives said transmitted trace information at said first data processing node; and in response to receipt of said trace

information, said first interrupt handling means causes a process-private interrupt of said trace process; said first data processing node further including means for reading said trace information, from said first data exchange means in response to said process-private interrupt.

In other aspects, the present invention provides a computer program for remote tracing of data processing nodes in an asynchronous messaging network

In a further preferred aspect of the present invention, there is provided a method in which the step of sending a trace command further comprises the following steps. Firstly, a trace command is written from a trace process into the local data exchange means. In response to the trace command, the process private interrupt on the local data processing node is initiated. Next, the trace command is transmitted across the network and then the process-private interrupt is replicated on the remote data processing node, in response to the transmitted trace command. An end user can dynamically issue trace commands in order to alter the processes to be traced and the nature of the trace information returned. Therefore, the user has real-time control over the remote tracing operations. This has considerable advantages over the static tracing operations mentioned in the prior art, since diagnostic information from remote data processing nodes can be viewed "on-the-fly".

According to a preferred embodiment of the present invention, once a process private interrupt has been caused of the target process, the process private interrupt is re-enabled. Preferably once the target process has written trace information into the remote data exchange means, the target process is re-started from the beginning of its execution.

According to a preferred embodiment of the present invention, the remote and local data exchange means are mailboxes. Furthermore separate mailboxes for read and write operations respectively, can be implemented. Each pair of read/write mailboxes are connected and in a system with fewer components, this is advantageous because the input/output operations can be easily differentiated

Preferably, the trace information is encrypted on the remote data processing node prior to transmission, and is decrypted on the local data processing node upon receipt of the trace information. This is advantageous in a customer environment, since the trace information transmitted is

secure. Preferably, the trace information is annotated so that it is concise and the amount of information transmitted across the communications medium is limited, ensuring that there is efficient data exchange between the data processing nodes.

5

In a further preferred aspect of the present invention the network is the Internet, however there is no limitation on the type of communications medium between the remote and local data processing nodes. For example, SNA could just as well be used. However, using TCP/IP, which is the main Internet protocol, is beneficial since this is the most widespread method for connecting disparate data processing nodes together. Additionally, executing the remote tracing operation over large distances where the data processing nodes are not in close proximity to each other is trivial.

10

15

In other preferred aspects of the present invention, the trace information transmitted across the communications medium can be viewed with a monitor connected to the data processing nodes.

20

Thus the present invention can be used as a diagnostic tool in conjunction with debugging operations. Since remote tracing operations are less tightly coupled to the processes to be traced compared with debugging operations, the results from the tracing can be used as an indicator as to where a problem lies. Subsequent to remote tracing operations, debugging operations can be executed to analyse the isolated problem in more depth.

25

Brief Description of the Drawings

30

The present invention will now be described, by way of example only, with reference to preferred embodiments thereof as illustrated in the following drawings:

35

Figure 1 shows a schematic diagram of a prior art system of inter-process communication within a single data processing node;

Figure 2 is a block diagram of communicating network servers in which the present invention may be implemented;

40

Figure 3 shows a schematic diagram of a distributed data processing system for tracing operations in remote data processing nodes according to the present invention;

Figure 4 shows a schematic diagram of a distributed data processing system for tracing operations in remote data processing nodes according to an alternative embodiment of the present invention; and

Figure 5 is a flow chart showing the steps of a method for remote tracing of operations in remote data processing nodes, according to the present invention, in the systems of Figures 3 and 4.

Detailed description of the Invention

Figure 1 illustrates schematically a known method of inter-process communication within a single data processing node. A single data processing node, as described herein, consists of a server or a client workstation.

The mailbox (20) is a software feature of the OpenVMS operating system and is used to exchange data between processes. (OpenVMS is a trademark of Compaq Computer Corporation.) The mailbox (20) is a pseudo device, which can be read and written, thus functioning as a buffer. A process (30) issues an I/O request to a mailbox (20) by declaring that a named procedure within the process (30) is to be run when information is written into the mailbox (20).

In the OpenVMS operating system, the named procedure is a process-private interrupt, known as an asynchronous system trap (AST). The AST is a software interrupt occurring during the execution of machine instructions in a process, allowing the instructions to be resumed from the beginning of execution. The initiation of the AST is caused by a process which is to receive notification during an I/O request.

To enable the AST to provide notification to the process (30) that information has been inputted into the mailbox (20), the process (30) includes a function modifier in the I/O request, known as the write attention (WRTATTN) function modifier.

Therefore, when the input process (10) writes information into the mailbox (20), this causes the mailbox (20) to deliver the attention WRATTN AST by queuing it to the process (30) that initiated the original I/O request. This technique allows the process (30) to continue executing and

performing other tasks, without it having to poll the mailbox (20) or otherwise stall. Hence, an I/O request is not outstanding at all times.

5 The notification from the AST results in the process (30) initiating an AST with the function modifier READATTN. The READATTN AST prompts the process (30) to read the information from the mailbox (20), therefore completing the I/O request. It is possible for multiple processes to request an AST for the same mailbox. Therefore, the first process to issue a request will complete the transfer of information.

10 For a more complete understanding of the present invention, attention is directed to Figure 2, in which the principal data structures and functional components of a computer system in which the invention is implemented are shown. Figure 2 illustrates two servers (75, 100) in
15 communication with each other via a network (105). In a preferred embodiment of the present invention, the data processing nodes including the servers (75, 100) utilise the OpenVMS operating system (97). Preferably, the network (105) is the Internet.

20 In normal operation, the server (75) of a remote data processing node executes a process (50), which is the target for the remote tracing operation. In this example, the process (50) is executing as part of the asynchronous message queuing product (96) MQSeries, which is installed on the remote server (75). (MQSeries is a trademark of IBM corporation.) The
25 MQSeries client program can also be installed on client workstations connected to the remote server (75). With the MQSeries product, there are different types of processes running, sometimes existing in multiple copies. Any of these processes are eligible as a target for the remote tracing operation.

30 A remote network server process (60) communicates with the target process (50) utilising a data exchange means (55) and an interrupt handling means (70), which are features of the operating system (97). Also located on the remote server (75), is an existing local trace facility (65), which
35 can be used within the single data processing node (75) to execute static trace commands.

40 A trace process (80), which initiates the remote tracing operation, is executing on a local server (100). A local network server process (90) communicates with the trace process (80) utilising a data exchange means (85) and an interrupt handling means (95). The local network server process

(90) and the remote network server process (60) communicate with each other using the network (105).

Although the present invention has been applied to asynchronous message queuing software (96), it should be understood that the present invention could be applied to any other type of process or software program.

Figure 3 illustrates schematically a tracing operation between remote data processing nodes. Figure 3 is used in conjunction with Figure 5, to provide a detailed description of the operations of Figure 2. Mailbox (55) is located on the remote server (75) and mailbox (85) is located on the local server (100). Every process, be it target process (50) or trace process (80), has a mailbox of its own. For example, during a one way I/O request, a trace process will input trace commands into the mailbox of the target process. It will complete the I/O request by reading the resulting output from its own mailbox.

As shown in Figure 2, the remote network server process (60) and the local network server process (90) communicate via the network (105). According to the preferred embodiment of the present invention, TCP/IP is utilised as the communications medium. Alternatively DECnet for OpenVMS or SNA could also be used. (DECnet is a trademark of Compaq Computer Corporation.)

Referring to Figure 5 illustrated in flow diagram form is the tracing operation between remote data processing nodes. The trace process (80) creates a mailbox (55) temporarily on the remote data processing node and the present invention enables the process (80) to write (500) a trace command into the mailbox (55) via the network (105). The target process (50) can distinguish between trace processes (80) and general processes, since the trace (80) and general processes utilise separate mailboxes. However, the use of separate mailboxes may cause a processing overhead, hence, alternatively, the trace processes (80) may be tagged with an identifier.

Trace commands typically report on function exit and entry points encountered during execution of the process being traced. It is possible to specify the range of functions being reported. Specifically, the target process (50) collects a certain amount of trace material automatically.

Specifically, an example of the use of trace commands in MQSeries is described. For instance, a record is kept on a stack of all the function calls that are currently active.

E.g. f called g called h

Therefore, h is currently executing. A trace program called monmq has a "show stack" command. The following trace commands are typed:

```
MQT> connect pid = 0x132
MQT> show stack
MQT> disconnect pid = 0x132
```

The first command establishes a connection to the process with id number 132 (hex). The second command uses trace mechanisms to obtain a snapshot of the current function call stack. The third command disconnects the process with id number 132 (hex).

Preferably, the trace process (80) can dynamically select any subset of the target processes to monitor, by sending further trace commands to the mailboxes associated with the selected subset and not to any other mailbox. The trace commands can also be dynamically varied to specify the nature of the trace material returned.

Once the trace command is written (500) into the mailbox (55), it is copied to a message block in a nonpaged pool, where it is stored until it is read. Since information in the form of a trace command has been written (500), an AST with the function modifier WRTATTN, is queued (505) to the target process (50), providing notification of the information.

The notification prompts the target process (50) to issue (505) a READATTN AST. This enables the process (50) to read (510) the trace command from the mailbox (55) and execute it. The present invention allows the target process (50) to respond to the command and enables the target process (50) to collect information regarding its own activity. The process (50) then writes (515) the trace information of its activity into the mailbox (55) of the remote data processing node. Preferably, the trace information is encrypted for added security.

In the present invention, step 515 initiates the trigger (520) of the interrupt, WRTATTN AST, notifying a process (60) running on the remote server that information is available in the mailbox (55). By issuing (520) a READATTN AST, the remote network server process (60) reads (530) the trace information from the mailbox (55). The remote network server process

(60) then communicates with the local network server process (90) over the network (105), transmitting (535) the trace information to the local network server process (90).

5 The code in the trace information is preferably annotated so that it is concise and consists of for example, a few dozen bytes. The annotation limits the amount of information being transmitted and therefore improves the performance of the remote tracing operation. Additionally the efficiency of the remote tracing operation can also be improved by running
10 the operation out of office hours and when the network bandwidth is satisfactory. Thus modest demands are made of the round trip time and the network. Utilising these measures will ensure that the results from the remote tracing operation will be received in an acceptable time frame and also that the operation does not have a substantial overhead associated
15 with it.

 The present invention allows the local network server process (90) to receive (540) the trace information, which it preferably decrypts. The local network server process (90) writes (545) the trace information into
20 the mailbox (85) of the local data processing node. The events occurring on the remote data processing node have now been replicated on the local data processing node. This replication process differs from prior art such as Lotus Notes where long-term copies are made and with other techniques such as disk mirroring. The remote tracing operation is more efficient, since no
25 lasting records are held. The replication of the environment, in that trace information has been written (545) into the local mailbox (85), causes the **WRTATTN AST** to be queued (550) to the trace process (80). In the present invention, the subsequent trigger (550) of a **READATTN AST** by the trace process (80) prompts the process (80) to read (555) the trace information
30 from the mailbox (85).

 Next, the end user decides (560) on the format of the output of the trace information. The trace information is either viewed (562) with a monitor of the local data processing node if the end user requires
35 immediate results, or it is written (564) into a trace file for consequent perusal.

 Next, the remote network server process (60) re-arms (570) the **AST** interrupt, if the operating system call made by the remote network server process (60) in order to read trace information from the mailbox (55) is to
40

be repeated. The re-enabling of the ASTs is required since all the data is removed after each delivery of the AST.

5 If no further trace information is to be read, the end user decides (575) whether the tracing operation is to continue. If further tracing is required, the tracing operation is repeated when the trace process (80) writes (500) trace commands into the mailbox (55). If no further tracing is required, the trace process (80) is terminated (580). The termination occurs when an end user completes issuing further trace commands from the local data processing node. Alternatively the termination process could be 10 automated and this is beneficial if for example, the tracing operation is to be completed after a certain period of time, or if the operation is to be executed during off-peak hours.

15 Figure 4 shows an alternative embodiment of the system as illustrated in Figure 3. Shown are separate mailboxes for input and output operations, namely, write and read operations, respectively. Specifically, the input (57, 59) and output (56, 58) mailboxes are connected in order to exchange data between them.

20 Implementing separate mailboxes is simpler than implementing a shared mailbox, since input and output operations are easily differentiated. However, prompting a process to read from one mailbox and write into a separate mailbox may become complicated especially as the number of processes increases. Moreover, a higher overhead is associated with 25 implementing multiple mailboxes. The implementation of a single mailbox for all input and output operations is more complex as the operations need to be distinguished. However, it is easier to use since all operations can occur via only one mailbox.

30 In a further preferred embodiment of the present invention, the remote tracing operation can also be applied to unresponsive or "hung" target processes. An unresponsive process may have terminated abnormally resulting in the suspension of any tasks the process was responsible for. Since the process is unresponsive to any input, determination of the cause of 35 termination is not possible. One solution is to reboot the data processing node, although this may result in data loss.

40 Specifically, it is an OpenVMS feature for an interrupt to execute at a higher priority level than a process. Therefore, even though a process may be unresponsive, an interrupt can continue to execute, sending trace

information detailing the possible cause of abnormal termination to the local data processing node. In this way, some degree of remote enquiry into the state of unresponsive, stalled target processes can be made.

5 Beneficially, by using a process-private interrupt, the trace process (80) can continue executing, since the interrupt will automatically execute once information is available in the mailbox (85). Therefore, the trace process (80) does not have to poll the mailbox (85) or otherwise stall. Alternatively, if the trace process (80) issues a simple read request from
10 the mailbox (85), it will have to wait until the target process (50) has completed the write operation into the mailbox (85) before it can complete the I/O request. Therefore, the trace process (80) cannot perform any additional processing until the I/O request is completed.

15 By utilising asynchronous ASTs this problem is overcome. For example, on the local data processing unit, the end user may be involved in other tasks, such as, browsing through previous trace material. The end user can continue browsing without interruption, until information is written into the mailbox. Also, more importantly, in a customer environment, the
20 intrusion into the executing processes on the remote data processing node is minimised, since the processes targeted can continue to do useful work until trace commands are written into the mailbox.

25 The dynamic control over the tracing operation and the limited amount of data being transmitted across a network ensures the operation is flexible and also improves the response times. Additionally, the present invention can be integrated into existing solutions easily, since standard features of an operating system, such as, data exchange means and interrupt handling means are utilised. Furthermore, the preferred communications
30 medium, TCP/IP, is a common protocol for connecting distributed systems. This further facilitates integration into existing systems.

35 Generally, the advantages of remote tracing, namely the ability to perform diagnosis of networked data processing nodes, coupled with the benefits of implementing mailboxes and process-private interrupts, provides a less intrusive method of problem determination over a network, whilst also improving the performance times.

CLAIMS

1. A method for remote tracing from a local data processing node of the execution of a process within an application program running on a remote data processing node in a distributed data processing network, said application program including its own local trace facility, said nodes communicating by asynchronous messaging via a data exchange means and each node including process-private interrupt handling means for indicating the presence of a command for the respective process in said data exchange means, said method comprising the steps of:

 sending a trace command from a trace process running on said local data processing node into a data exchange means of said remote data processing node;

 in response to said trace command, causing a process-private interrupt of a target process running on said remote data processing node;

 in response to said process-private interrupt, said target process writing trace information from said trace facility to said data exchange means;

 transmitting said trace information across said network;

 receiving in a data exchange means on said local data processing node, said transmitted trace information;

 in response to receiving said trace information, causing a process-private interrupt of said trace process; and

 in response to said process-private interrupt, reading said trace information by said trace process, from said local data exchange means.

2. A method as claimed in claim 1, in which the step of sending a trace command further comprises the steps of:

 writing a trace command from a trace process into said local data exchange means;

 initiating said process private interrupt on said local data processing node in response to said trace command;

transmitting said trace command across the network; and

5 replicating said process-private interrupt on said remote data processing node in response to said transmitted trace command.

10 3. A method as claimed in claim 1, in which after the step of causing a process private interrupt of said target process, said process private interrupt is re-enabled.

4. A method as claimed in claim 1, in which after said step of writing trace information, said target process is re-started from the beginning of its execution.

15 5. A method as claimed in any preceding claim, in which said remote and local data exchange means are mailboxes.

20 6. A method as claimed in claim 5 in which each of said data exchange means comprise separate mailboxes for said read and said write operations, respectively.

7. A method as claimed in any preceding claim, in which said trace information is encrypted on said remote data processing node.

25 8. A method as claimed in claim 7, in which said trace information is decrypted on said local data processing node.

30 9. A method as claimed in any preceding claim, in which said trace information is annotated.

10. A method as claimed in any preceding claim in which said local and remote data processing nodes communicate over the Internet.

35 11. A method as claimed in any preceding claim in which said trace information is viewed with a monitor of said data processing nodes.

12. A computer program carrying out the method of any of claims 1 to 11.

40 13. A distributed data processing system comprising a plurality of data processing nodes connected via a network, each node having a processor, memory and operating system capable of executing application programs, each

of said operating systems including data exchange means and interrupt handling means, a first of said nodes comprising:

5 means for sending a trace command from a trace process running on said first data processing node into the data exchange means of a second data processing node, said second data processing node including a trace facility for tracing the execution of a process within an application program running on said second node,

10 said interrupt handling means of said second data processing node, in response to said trace command, causing a process-private interrupt of a target process running on said second data processing node;

15 said second data processing node further including means for writing trace information from the trace facility, to said second data exchange means in response to the process-private interrupt; and

20 means for transmitting said trace information across said network whereby said first data exchange means receives said transmitted trace information at said first data processing node; and

in response to receipt of said trace information, said first interrupt handling means causes a process-private interrupt of said trace process;

25 said first data processing node further including means for reading said trace information, from said first data exchange means in response to said process-private interrupt.

ABSTRACT

REMOTE TRACING OF DATA PROCESSING NODES
IN AN ASYNCHRONOUS MESSAGING NETWORK

5
Remote tracing from a local data processing node of the execution of a process within an application program running on a remote data processing node in a distributed data processing network. The application program
10 includes its own local trace facility and the nodes communicate by asynchronous messaging via mailboxes. Each node includes process-private interrupt handling means for indicating the presence of a command for the respective process in the mailboxes.

15 A trace command is sent from a trace process running on a local data processing node into a mailbox of a remote data processing node. In response to the trace command, a process-private interrupt is caused of a target process running on the remote data processing node. In response to the interrupt, the target process writes trace information from the trace
20 facility to the mailbox. This trace information is transmitted across the network and is received in a mailbox on the local data processing node. In response to receiving the trace information, a process-private interrupt is caused of the trace process. In response to the process-private interrupt the trace information is read by the trace process from the mailbox.

25

30


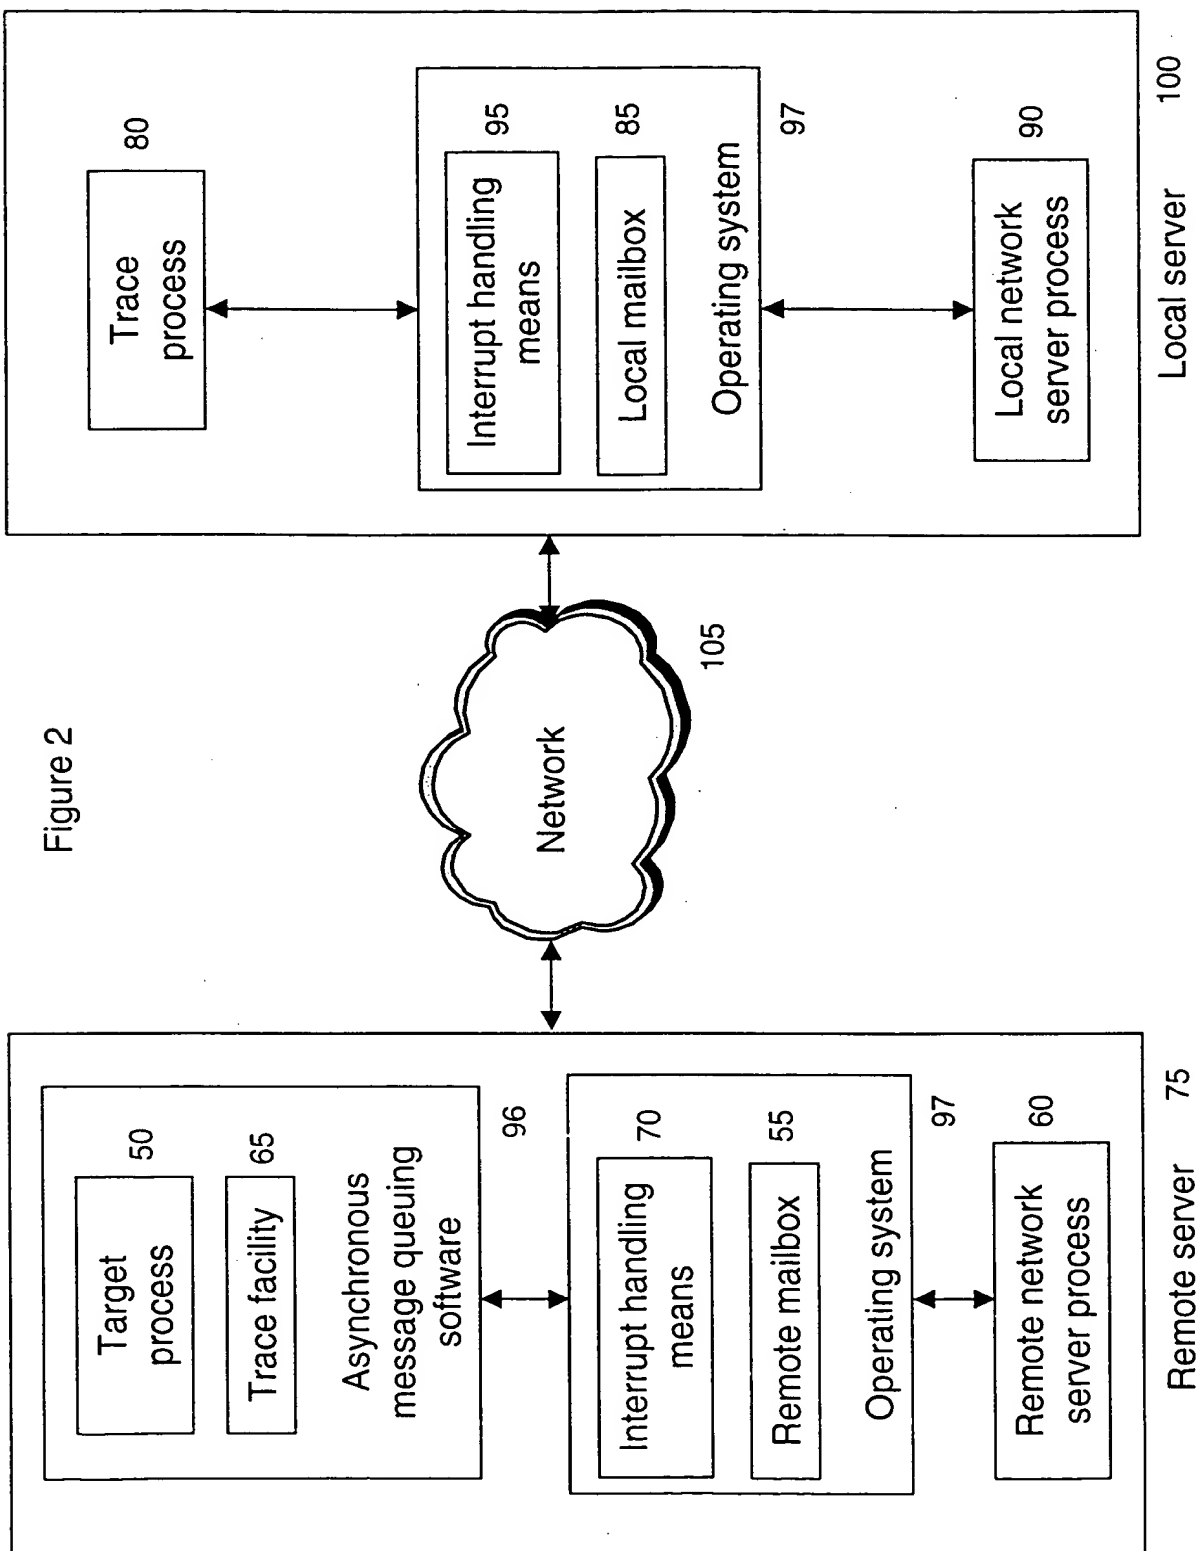
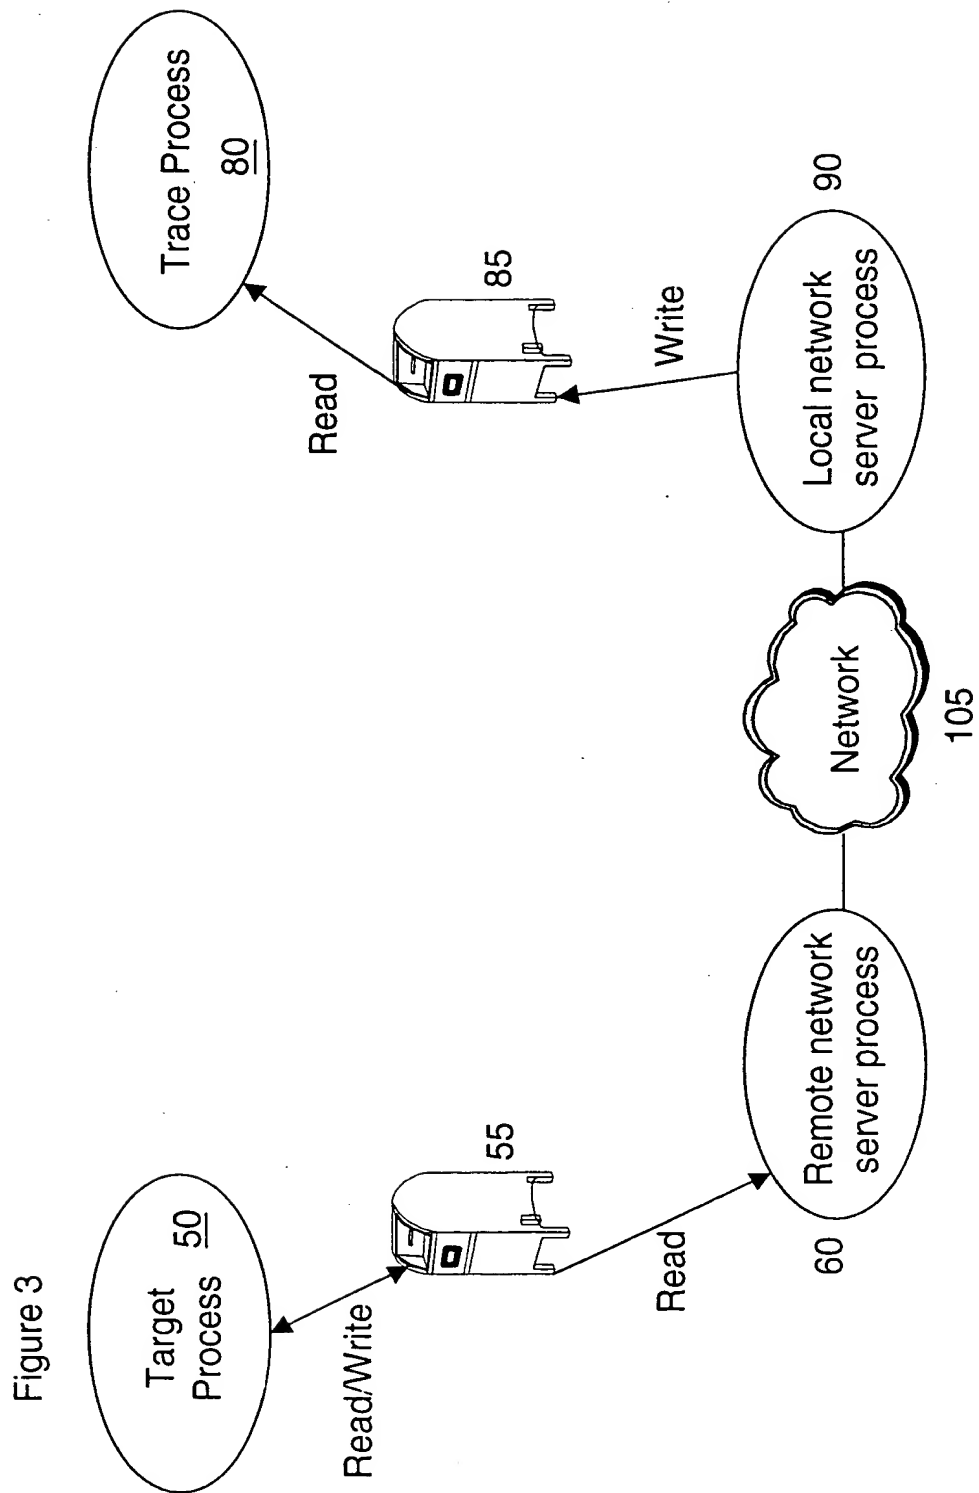


Figure 1
Prior Art





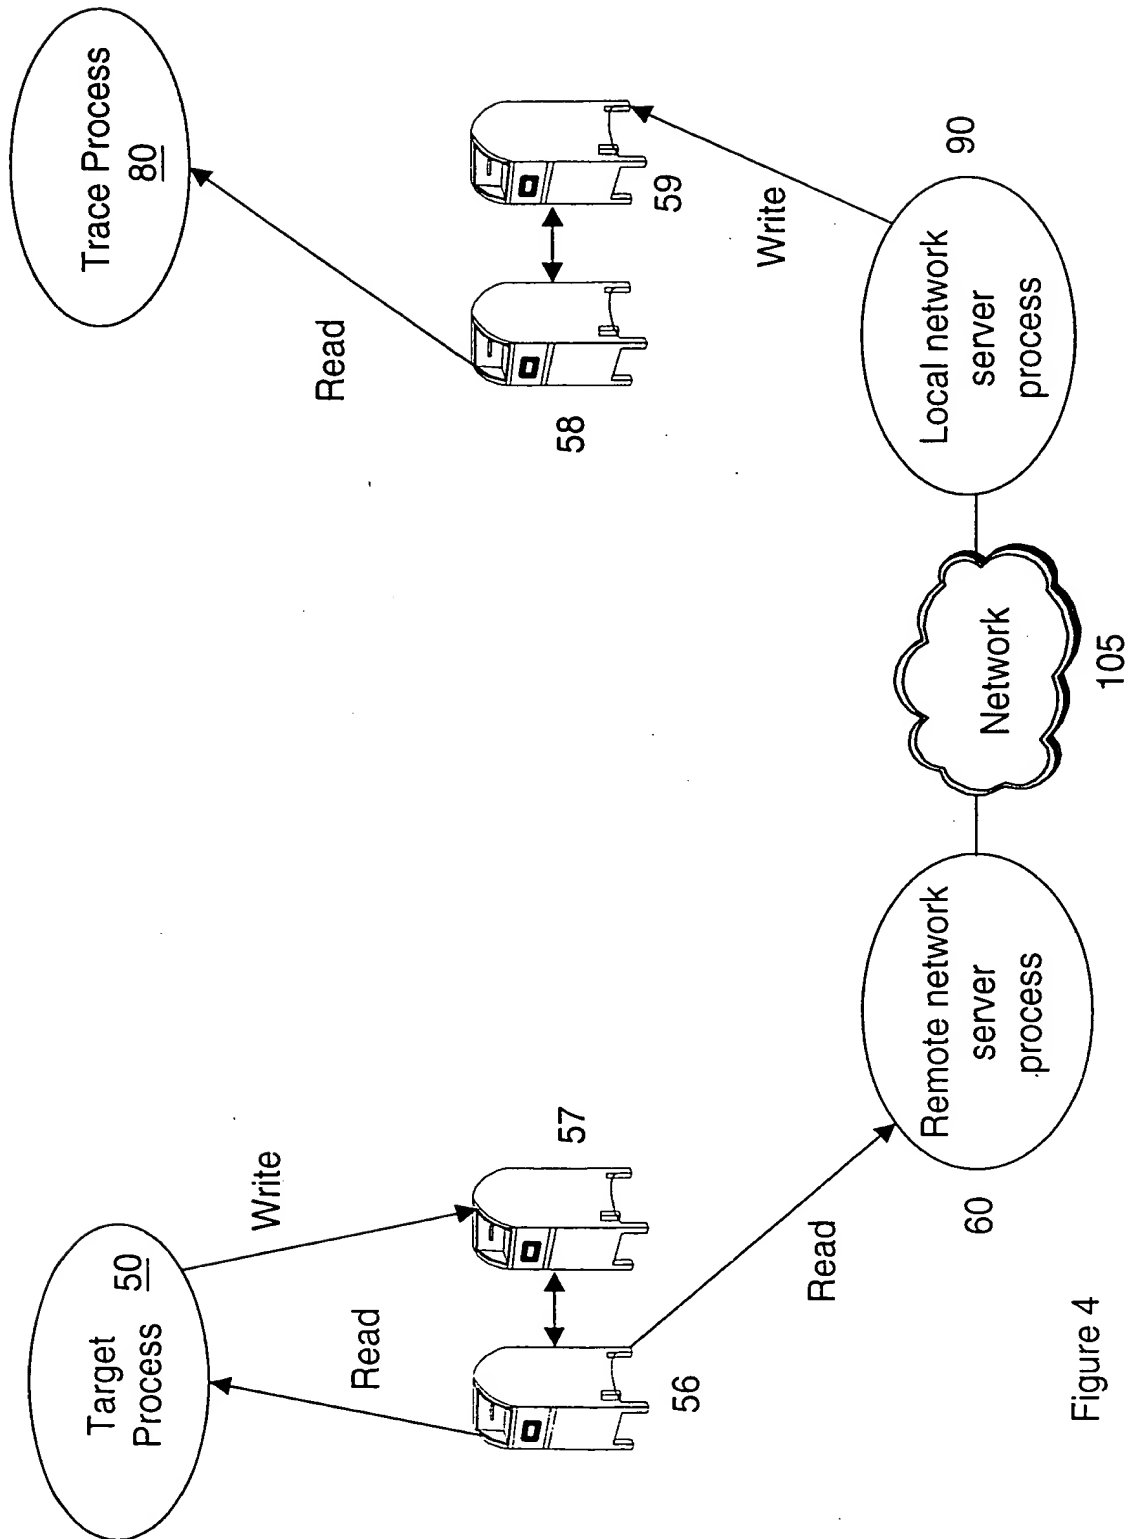


Figure 4

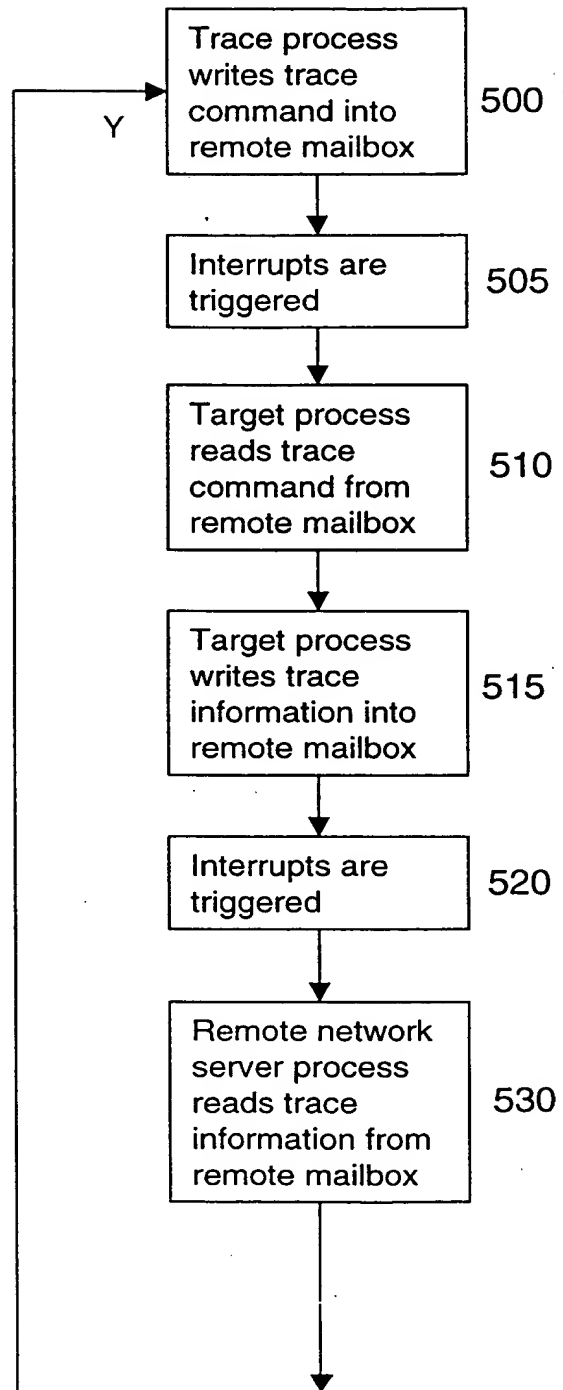
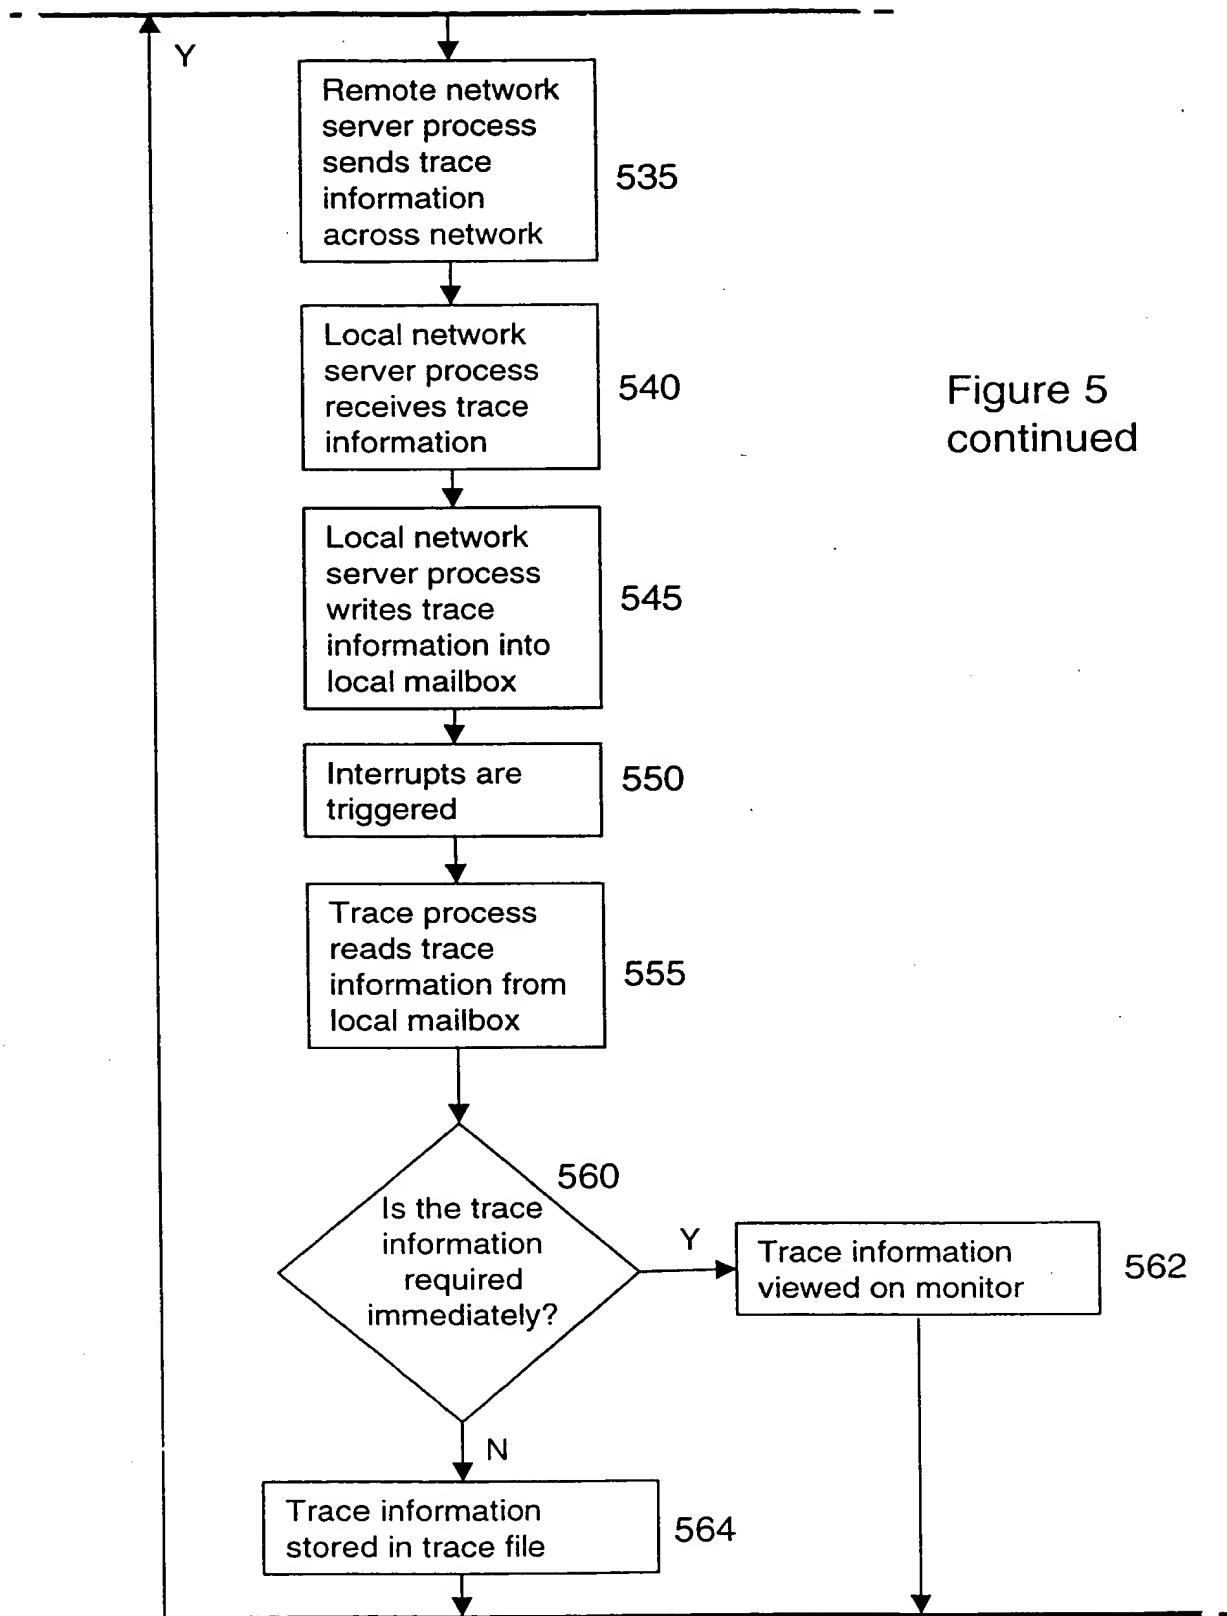


Figure 5



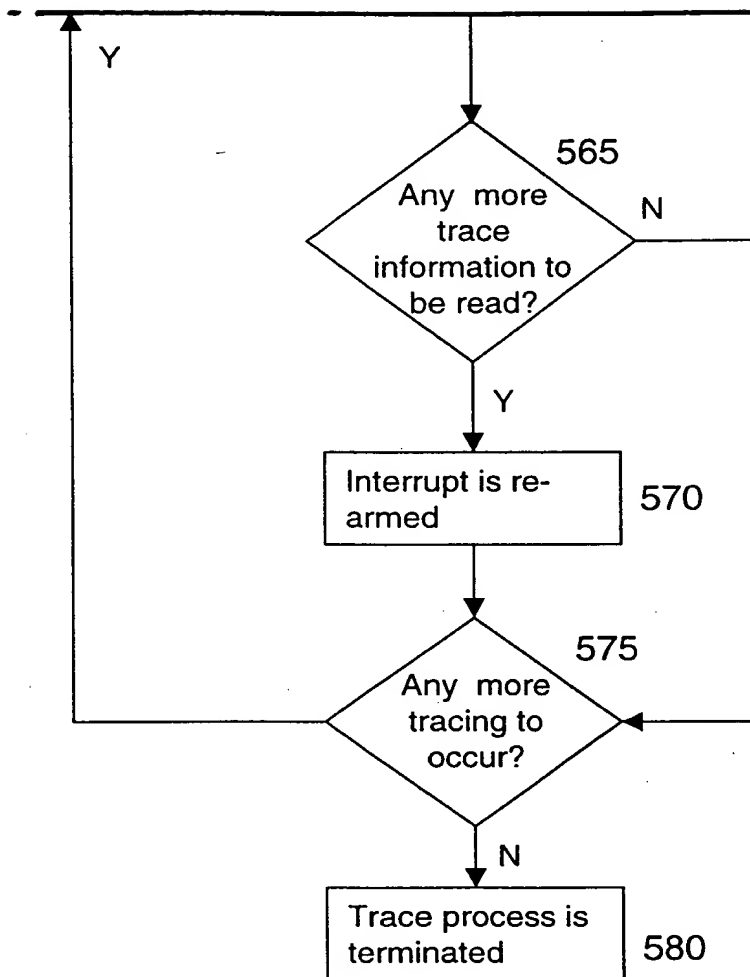


Figure 5
continued

THIS PAGE BLANK (USPTO)